

On the Inductive Bias of Masked Language Modeling: From Statistical to Syntactic Dependencies

Tianyi Zhang

Computer Science Department
Stanford University

tianyizhang@cs.stanford.edu

Tatsunori B. Hashimoto

Computer Science Department
Stanford University

thashim@stanford.edu

Abstract

We study how masking and predicting tokens in an unsupervised fashion can give rise to linguistic structures and downstream performance gains. Recent theories have suggested that pretrained language models acquire useful inductive biases through masks that implicitly act as cloze reductions. While appealing, we show that the success of the random masking strategy used in practice cannot be explained by such cloze-like masks alone. We construct cloze-like masks using task-specific lexicons for three different classification datasets and show that the majority of pretrained performance gains come from *generic* masks that are not associated with the lexicon. To explain the empirical success of these generic masks, we demonstrate a correspondence between the masked language model (MLM) objective and existing methods for learning statistical dependencies in graphical models. Using this, we derive a method for extracting these learned statistical dependencies in MLMs and show that these dependencies encode useful inductive biases in the form of syntactic structures. In an unsupervised parsing evaluation, simply forming a minimum spanning tree on the implied statistical dependence structure outperforms a classic method for unsupervised parsing (58.74 vs. 55.91 UAS).

1 Introduction

Pretrained masked language models (Devlin et al., 2019; Liu et al., 2019b) have benefitted a wide range of natural language processing (NLP) tasks (Liu, 2019; Wadden et al., 2019; Zhu et al., 2020). Despite recent progress in understanding *what* useful information is captured by MLMs (Liu et al., 2019a; Hewitt and Manning, 2019), it remains a mystery *why* task-agnostic masking of words can capture linguistic structures and transfer to downstream tasks.

One popular justification of MLMs relies on viewing masking as a form of cloze reduction.




Figure 1: We study the inductive bias of MLM objectives and show that cloze-like masking (left) does not account for much of the downstream performance gains. Instead, we show that MLM objectives are biased towards extracting both statistical and syntactic dependencies using random masks (right).

Cloze reductions reformulate an NLP task into a prompt question and a blank and elicit answers by filling in the blank (Figure 1). When tested by cloze reductions pretrained MLMs and left-to-right language models (LMs) have been shown to possess abundant factual knowledge (Petroni et al., 2019) and display impressive few-shot ability (Brown et al., 2020). This success has inspired recent hypotheses that some word masks are cloze-like and provide indirect supervision to downstream tasks (Saunshi et al., 2020; Lee et al., 2020). For example, a sentiment classification task (Pang et al., 2002) can be reformulated into filling in *like* or *hate* in the cloze *I [MASK] this movie*. Such cloze-like masks provide a clear way in which an MLM can implicitly learn to perform sentiment classification.

While this hypothesis is appealing, MLMs in practice are trained with uniform masking that does not contain the special structure required by cloze-like masks most of the time. For example, predicting a generic word *this* in the cloze *I like [MASK] movie* would not offer task-specific supervision. We quantify the importance of cloze-like and generic masks by explicitly creating cloze-like masks using task-specific lexicons and comparing models pretrained on these masks. These experiments suggest that although cloze-like masks can be helpful, the success of uniform masking cannot

be explained via cloze-like masks alone. In fact, we demonstrate that uniform masking performs as well as a negative control where we explicitly remove cloze-like masks from the mask distribution.

To address this mismatch between theory and practice, we offer a new hypothesis of how generic masks can help downstream learning. We propose a conceptual model for MLMs by drawing a correspondence between masking and graphical model neighborhood selection (Meinshausen and Bühlmann, 2006). Using this, we show that MLM objectives are designed to recover statistical dependencies in the presence of latent variables and propose an estimator that can recover these learned dependencies from MLMs. We hypothesize that statistical dependencies in the MLM objective capture useful linguistic dependencies and demonstrate this by using recovered statistical dependencies to perform unsupervised parsing, outperforming an actual unsupervised parsing baseline (58.74 vs 55.91 UUAS; Klein and Manning, 2004). We release our implementation on Github¹.

2 Related works

Theories inspired by Cloze Reductions. Cloze reductions are fill-in-the-blank tests that reformulate an NLP task into an LM problem. Existing work demonstrates that such reductions can be highly effective for zero/few-shot prediction (Radford et al., 2019; Brown et al., 2020) as well as relation extraction (Petroni et al., 2019; Jiang et al., 2020).

These fill-in-the-blank tasks provide a clear way by which LMs can obtain supervision about downstream tasks, and recent work demonstrates how such implicit supervision can lead to useful representations (Saunshi et al., 2020). More general arguments by Lee et al. (2020) show these theories hold across a range of self-supervised settings. While these theories provide compelling arguments for the value of pre-training with cloze tasks, they do not provide a clear reason why *uniformly random* masks such as those used in BERT provide such strong gains. In our work, we quantify this gap using lexicon-based cloze-like masks and show that cloze-like masks alone are unlikely to account for the complete success of MLM since generic and non-cloze masks are responsible for a substantial part of the empirical performance of MLMs.

¹https://github.com/tatsu-lab/mlm_inductive_bias

Theories for vector representations. Our goal of understanding how masking can lead to useful inductive biases and linguistic structures is closely related to that of papers studying the theory of word embedding representations (Mikolov et al., 2013; Pennington et al., 2014; Arora et al., 2015). Existing work has drawn a correspondence between word embeddings and low-rank factorization of a pointwise mutual information (PMI) matrix (Levy and Goldberg, 2014) and others have shown that PMI is highly correlated with human semantic similarity judgements (Hashimoto et al., 2016).

While existing theories for word embeddings cannot be applied to MLMs, we draw inspiration from them and derive an analogous set of results. Our work shows a correspondence between MLM objectives and graphical model learning through conditional mutual information, as well as evidence that the conditional independence structure learned by MLMs is closely related to syntactic structure.

Probing Pretrained Representations. Recent work has applied probing methods (Belinkov and Glass, 2019) to analyze *what* information is captured in the pretrained representations. This line of work shows that pretrained representations encode a diverse range of knowledge (Peters et al., 2018; Tenney et al., 2019; Liu et al., 2019a; Hewitt and Manning, 2019; Wu et al., 2020). While probing provides intriguing evidence of linguistic structures encoded by MLMs, they do not address the goals of this work, which is *how* the pretraining objective encourages MLMs to extract such structures.

3 Motivation

3.1 Problem Statement

Masked Language Modeling asks the model to predict a token given its surrounding context. Formally, consider an input sequence X of L tokens $\langle x_1, \dots, x_L \rangle$ where each variable takes a value from a vocabulary \mathcal{V} . Let $X \sim D$ be the data generating distribution of X . Let x_i be the i th token in X , and let $X_{\setminus i}$ denote the sequence after replacing the i th token with a special [MASK] token. In other words,

$$X_{\setminus i} := \langle x_1, \dots, x_{i-1}, [\text{MASK}], x_{i+1}, \dots, x_L \rangle.$$

Similarly, define $X_{\setminus \{i,j\}}$ as replacing both x_i and x_j with [MASK]. MLM determines what tokens are masked by a mask distribution $i \sim M$. The goal of MLM is to learn a probabilistic model p_θ

Modified Input:	<i>beautiful movie</i> <i>positive</i>
Cloze-like Mask:	<i>beautiful movie</i> [MASK]
Generic Mask:	<i>beautiful</i> [MASK] <i>positive</i>

Figure 2: In our case study, we append the **true label** to each input and create ideal **cloze-like masks**. We study how deviations from the ideal mask distribution affect downstream performance by adding in **generic masks**.

that minimizes

$$L_{\text{MLM}} = \mathbb{E}_{X \sim D, i \sim M} -\log p_{\theta}(x_i | X_{\setminus i}).$$

In BERT pretraining, each input token is masked with a fixed, uniform probability, which is a hyperparameter to be chosen. We refer to this strategy as *uniform masking*.

Finetuning is the canonical method for using pretrained MLMs. Consider a prediction task where $y \in \mathcal{Y}$ is the target variable, *e.g.*, the sentiment label of a review. Finetuning uses gradient descent to modify the pretrained parameters θ and learn a new set of parameters ϕ to minimize

$$L_{\text{finetune}} = \mathbb{E}_{X \sim D', y \sim p(y|X)} -\log p_{\theta, \phi}(y|X),$$

where $p(y|x)$ is the ground-truth distribution and D' is the data distribution of the downstream task.

Our goals. We will study how the mask distribution M affects downstream performance. We define perfect cloze reductions as some partition of the vocabulary \mathcal{V}_y such that $p(x_i \in \mathcal{V}_y | X_{\setminus i}) \approx p(y|X)$. For a distribution M such that the masks we draw are perfect cloze-reductions, the MLM objective offers direct supervision to finetuning since $L_{\text{MLM}} \approx L_{\text{finetune}}$. In contrast to cloze-like masking, in uniform masking we can think of p_{θ} as implicitly learning a generative model of \mathbf{X} (Wang and Cho, 2019). Therefore, as M moves away from the ideal distribution and becomes more uniform, we expect p_{θ} to model more of the full data distribution D instead of focusing on cloze-like supervision for the downstream task. This mismatch between theory and practice raises questions about how MLM with uniform masking can learn useful inductive biases.

When L_{MLM} is not L_{finetune} , what is L_{MLM} learning? We analyze L_{MLM} and show that it is similar to a form of conditional mutual information based graphical model structure learning.

3.2 Case Study for Cloze-like Masking

To motivate our subsequent discussions, we perform a controlled study for the case when $L_{\text{MLM}} \approx$

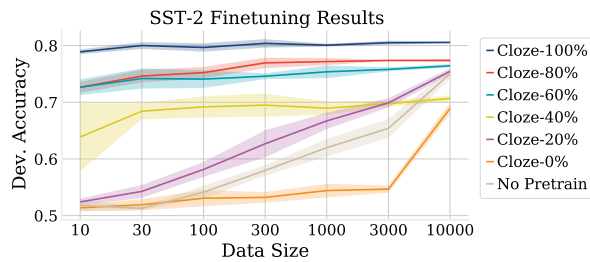


Figure 3: SST-2 development set accuracy. CLOZE- $p\%$ is pretrained on a mixture of masks where $p\%$ of the masks are Cloze-like. NOPRETRAIN trains a classifier without any pretraining. Even a small modification of the ideal mask distribution degrades performance.

L_{finetune} and analyze how deviations from the ideal mask distribution affect downstream performance. We perform analysis on the Stanford Sentiment Treebank (SST-2; Socher et al., 2013), which requires models to classify short movie reviews into positive or negative sentiment. We append the ground-truth label (as the word *positive* or *negative*) to each movie review (Figure 2). Masking the last word in each review is, by definition, an ideal mask distribution. To study how the deviation from the ideal mask distribution degrades downstream performance, we vary the amount of cloze-like masks during training. We do this by masking out the last word for $p\%$ of the time and masking out a random word in the movie review for $(100 - p)\%$ of the time, and choose $p \in \{0, 20, 40, 60, 80, 100\}$.

Experimental details. We split the SST-2 training set into two halves, use one for pretraining, and the other for finetuning. For the finetuning data, we do not append the ground-truth label. We pre-train small transformers with L_{MLM} using different masking strategies and finetune them along with a baseline that is not pretrained (NOPRETRAIN). Further details are in Appendix A.

Results. We observe that while cloze-like masks can lead to successful transfer, *even a small modification of the ideal mask distribution deteriorates performance*. Figure 3 shows the development set accuracy of seven model variants averaged across ten random trials. We observe as p decreases, the performance of CLOZE- $p\%$ degrades. Notably, CLOZE-80% is already worse than CLOZE-100% and CLOZE-20% does not outperform NOPRETRAIN by much. We notice that CLOZE-0% in fact degrades finetuning performance, potentially because the pretrained model is over-specialized to the language modeling task (Zhang et al., 2020; Tamkin et al., 2020). While this is a toy example, we observe similar results for actual MLM models




Figure 4: Our conceptual framework of MLM. All coordinates of X are dependent on the latent variable Z while there is only sparse dependency among X .

across three tasks (Section 5.1), and this motivates us to look for a framework that explains the success of generic masks in practice.

4 Analysis

In the previous section, we saw that cloze-like masks do not necessarily explain the empirical success of MLMs with uniform masking strategies. Understanding uniform masking seems challenging at first, as uniform-mask MLMs seem to lack task-specific supervision and is distinct from existing unsupervised learning methods such as word embeddings (which rely upon linear dimensionality reduction) and autoencoders (which rely upon denoising). However, we show in this section that there is a correspondence between MLM objectives and classic methods for graphical model structure learning. As a consequence, we demonstrate that MLMs are implicitly trained to recover statistical dependencies among observed tokens.

4.1 Intuition and Theoretical Analysis

Our starting point is the observation that predicting a single feature (x_i) from all others ($X_{\setminus i}$) is the core subroutine in the classic Gaussian graphical model structure learning algorithm of Meinshausen and Bühlmann (2006). In this approach, L different Lasso regression models are trained (Tibshirani, 1996) with each model predicting x_i from $X_{\setminus i}$, and the nonzero coefficients of this regression correspond to the conditional dependence structure of the graphical model.

The MLM objective can be interpreted as a nonlinear extension of this approach, much like a classical algorithm that uses conditional mutual information (MI) estimators to recover a graphical model (Anandkumar et al., 2012). Despite the similarity, real world texts are better viewed as models with latent variables (e.g. topics; Blei et al., 2003) and many dependencies across tokens arise due to latent variables, which makes learning the direct dependencies difficult. We show that MLMs

implicitly recover the latent variables and can capture the direct dependencies while accounting for the effect of latent variables. Finally, MLMs are only approximations to the true distribution and we show that the MLM objective can induce high-quality approximations of conditional MI.

Analysis setup. To better understand MLMs as a way to recover graphical model structures, we show mask-based models can recover latent variables and the direct dependencies among variables in the Gaussian graphical model setting of Meinshausen and Bühlmann (2006). Let $\mathbf{X} = [x_1, \dots, x_L] \in \mathbb{R}^L$ represent an input sequence where each of its coordinates x_i represents a token, and $\mathbf{Z} \in \mathbb{R}^k$ be a latent variable that controls the sequence generation process. We assume that all coordinates of \mathbf{X} are dependent on the latent variable \mathbf{Z} , and there are sparse dependencies among the observed variables (Figure 4). In other words, we can write $\mathbf{Z} \sim N(0, \Sigma_{\mathbf{Z}\mathbf{Z}})$ and $\mathbf{X} \sim N(A\mathbf{Z}, \Sigma_{\mathbf{X}\mathbf{X}})$. Intuitively, we can imagine that \mathbf{Z} represents shared semantic information, e.g. a topic, and $\Sigma_{\mathbf{X}\mathbf{X}}$ represents the syntactic dependencies. In this Gaussian graphical model, the MLM is analogous to regressing each coordinate of \mathbf{X} from all other coordinates, which we refer to as masked regression.

MLM representations can recover latent variable. We now study the behavior of masked regression through the representation $\mathbf{x}_{\text{mask},i}$ that is obtained by applying masked regression on the i th coordinate of \mathbf{X} and using the predicted values. Our result shows that masked regression is similar to the two-step process of first recovering the latent variable \mathbf{Z} from $\mathbf{X}_{\setminus i}$ and then predicting x_i from \mathbf{Z} .

Let $\Sigma_{\mathbf{X}\mathbf{X},\setminus i,i} \in \mathbb{R}^{d-1}$ be the vector formed by dropping the i th row and taking the i th column of $\Sigma_{\mathbf{X}\mathbf{X}}$ and $\beta_{2\text{SLS},i}$ be the linear map resulting from the two-stage regression $\mathbf{X}_{\setminus i} \rightarrow \mathbf{Z} \rightarrow x_i$.

Proposition 1. Assuming that $\Sigma_{\mathbf{X}\mathbf{X}}$ is full rank,

$$\mathbf{x}_{\text{mask},i} = \beta_{2\text{SLS},i}\mathbf{X}_{\setminus i} + O(\|\Sigma_{\mathbf{X}\mathbf{X},\setminus i,i}\|_2),$$

In other words, masked regression implicitly recovers the subspace that we would get if we first explicitly recovered the latent variables ($\beta_{2\text{SLS},i}$) with an error term that scales with the off-diagonal terms in $\Sigma_{\mathbf{X}\mathbf{X}}$. The proof is presented in Appendix C.

To give additional context for this result, let us consider the behavior of a different representation learning algorithm: PCA. It is well-known that PCA can recover the latent variables as long as

the $\Sigma_{\mathbf{Z}\mathbf{Z}}$ dominates the covariance $\text{Cov}(\mathbf{X})$. We state this result in terms of \mathbf{X}_{PCA} , the observed data projected to the first k components of PCA.

Proposition 2. *Let λ_k be the k th eigenvalue of $A\Sigma_{\mathbf{Z}\mathbf{Z}}A^\top$ and $\lambda_{\mathbf{X}\mathbf{X},k+1}$ be the $k+1$ th eigenvalue of $\Sigma_{\mathbf{X}\mathbf{X}}$ and V be the first k eigenvectors of $\text{Cov}(\mathbf{X})$. Assuming $\lambda_k > \lambda_{\mathbf{X}\mathbf{X},k+1}$, we have*

$$\mathbb{E}_{\mathbf{X}} \|\mathbf{AZ} - \mathbf{X}_{\text{PCA}}\|_2 \leq \frac{\sqrt{2} \|\Sigma_{\mathbf{X}\mathbf{X}}\|_{op}}{\lambda_k - \lambda_{\mathbf{X}\mathbf{X},k+1}} (\|\mathbf{AZ}\|_2 + \sqrt{\text{tr}(\Sigma_{\mathbf{X}\mathbf{X}})}) + \|AA^\top\|_{op} \sqrt{\text{tr}(\Sigma_{\mathbf{X}\mathbf{X}})},$$

where $\|\cdot\|_{op}$ is the operator norm and $\text{tr}(\cdot)$ is the trace.

This shows that whenever $\Sigma_{\mathbf{X}\mathbf{X}}$ is sufficiently small and λ_k is large (i.e., the covariance is dominated by Z), then PCA recovers the latent information in \mathbf{Z} . The proof is based on the Davis-Kahan theorem (Stewart and Sun, 1990) and is presented in Appendix C.

Comparing the bound of PCA and masked regression, both bounds have errors that scales with $\Sigma_{\mathbf{X}\mathbf{X}}$, but the key difference in the error bound is that the error term for masked regression does not scale with the per-coordinate noise ($\text{diag}(\Sigma_{\mathbf{X}\mathbf{X}})$) and thus can be thought of as focusing exclusively on interactions within \mathbf{X} . Analyzing this more carefully, we find that $\Sigma_{\mathbf{X}\mathbf{X},\setminus i,i}$ corresponds to the statistical dependencies between x_i and $\mathbf{X}_{\setminus i}$, which we might hope captures useful, task-agnostic structures such as syntactic dependencies.

MLM log-probabilities can recover direct dependencies. Another effect of latent variables is that many tokens have indirect dependencies through the latent variables, which poses a challenge to recovering the direct dependencies among tokens. We now show that the MLMs can account for the effect of latent variable.

In the case where there are no latent variables, we can identify the direct dependencies via conditional MI (Anandkumar et al., 2012) because any x_i and x_j that are disconnected in the graphical model will have zero conditional MI, i.e., $I(x_i; x_j | X_{\setminus \{i,j\}}) = 0$. One valuable aspect of MLM is that we can identify direct dependencies even in the presence of latent variables.

If we naively measure statistical dependency by mutual information, the coordinates of \mathbf{X} would appear dependent on each other because they are all connected with \mathbf{Z} . However, the MLM objective resolves this issue by conditioning on $X_{\setminus \{i,j\}}$. We show that latent variables (such as topics) that are

easy to predict from $X_{\setminus \{i,j\}}$ can be ignored when considering conditional MI.

Proposition 3. *The gap between conditional MI with and without latent variables is bounded by the conditional entropy $H(\mathbf{Z} | X_{\setminus \{i,j\}})$,*

$$I(x_i; x_j | X_{\setminus \{i,j\}}) - I(x_i; x_j | \mathbf{Z}, X_{\setminus \{i,j\}}) \leq 2H(\mathbf{Z} | X_{\setminus \{i,j\}}).$$

This suggests that when the context $X_{\setminus \{i,j\}}$ captures enough of the latent information, conditional MI can remove the confounding effect of the shared topic \mathbf{Z} and extract the direct and sparse dependencies within \mathbf{X} (see Appendix C for the proof).

MLM objective encourages capturing conditional MI. We have now shown that conditional MI captures direct dependencies among tokens, even in the presence of latent variables. Next, we will show that the MLM objective ensures that a LM with low log-loss accurately captures the conditional MI. We now show that learning the MLM objective implies high-quality estimation of conditional MI. Denote $X(i, v)$ as substituting x_i with a new token v ,

$$X(i, v) = \langle x_1, \dots, x_{i-1}, v, x_{i+1}, \dots, x_L \rangle.$$

Conditional MI is defined as the expected pointwise mutual information (PMI) conditioned on the rest of the tokens,

$$I_p = \mathbb{E}_{x_i, x_j} [\log p(x_i | X_{\setminus i}(j, x_j)) - \log \mathbb{E}_{x_j | x_i} p(x_i | X_{\setminus i}(j, x_j))]]$$

where I_p is the abbreviation of $I_p(x_i; x_j | X_{\setminus \{i,j\}})$. Our main result is that the log-loss MLM objective directly bounds the gap between the true conditional mutual information from the data distribution and an estimator that uses the log-probabilities from the model. More formally,

Proposition 4. *Let*

$$\hat{I}_{p_\theta} = \mathbb{E}_{x_i, x_j} [\log p_\theta(x_i | X_{\setminus i}(j, x_j)) - \log \mathbb{E}_{x_j | x_i} p_\theta(x_i | X_{\setminus i}(j, x_j))]]$$

be an estimator constructed by the model distribution p_θ . Then we can show,

$$|\hat{I}_{p_\theta} - I_p| \leq \mathbb{E}_{x_j} D_{\text{kl}}(p(x_i | X_{\setminus i}(j, x_j)) \| p_\theta(x_i | X_{\setminus i}(j, x_j))),$$

where D_{kl} represents the KL-divergence.

Here, the KL-divergence corresponds to the L_{MLM} objective, up to a constant entropy term that depends on p . We present the proof in Appendix C. In other words, the MLM objective is implicitly encouraging the model to match its implied conditional MI to that of the data. We now use this result to create an estimator that extracts the conditional independence structures implied by MLM.

4.2 Extracting statistical dependencies implied by MLMs

Our earlier analysis in Proposition 4 suggests that an MLM with low loss has an accurate approximation of conditional mutual information. Using this result, we will now propose a procedure which estimates \hat{I}_{p_θ} . The definition of \hat{I}_{p_θ} shows that if we can access samples of x_i and x_j from the true distribution p , then we can directly estimate the conditional mutual information by using the log probabilities from the MLM. Unfortunately, we cannot draw new samples of $x_j | X_{\setminus\{i,j\}}$, leading us to approximate this distribution using Gibbs sampling on the MLM distribution.

Our Gibbs sampling procedure is similar to the one proposed in Wang and Cho (2019). We start with $X^0 = X_{\setminus\{i,j\}}$. For the t th iteration, we draw a sample x_i^t from $p_\theta(x_i | X_{\setminus i}^{t-1})$, and update by $X^t = X^{t-1}(i, x_i^t)$. Then, we draw a sample x_j^t from $p_\theta(x_j | X_{\setminus j}^t)$ and set $X^t = X^t(j, x_j^t)$. We repeat and use the samples $(x_i^1, x_j^1), \dots, (x_i^t, x_j^t)$ to compute the expectations for conditional MI.

This procedure relies upon an additional assumption that samples drawn from the MLM are faithful approximations of the data generating distribution. However, we show empirically that even this approximation is sufficient to test the hypothesis that the conditional independences learned by an MLM capture syntactic dependencies (Section 5.2).

5 Experiment

We now test two predictions from our analyses. First, similar to our observation in the case study, we show that cloze-like masks do not explain the success of uniform masks on three real-world datasets. Second, our alternative view of relating MLM to graphical models suggests that statistical dependencies learned by MLMs may capture linguistic structures useful for downstream tasks. We demonstrate this by showing that MLMs’ statistical dependencies reflect syntactic dependencies.

5.1 Uniform vs Cloze-like Masking

Setup. We now demonstrate that real-world tasks and MLMs show a gap between task-specific cloze masks and random masks. We compare the MLM with random masking to two different control groups. In the positive control (CLOZE), we pre-train with only cloze-like masks and in the negative control (NOCLOZE), we pre-train by explicitly excluding cloze-like masks. If the success of MLM

can be mostly explained by implicit cloze reductions, then we should expect CLOZE to have strong downstream performance while NOCLOZE leads to a minimal performance gain. We compare pre-training with the uniform masking strategy used in BERT (UNIFORM) to these two control groups. If UNIFORM performs worse than the positive control and more similar to the negative control, then we know that uniform masking does not leverage cloze-like masks effectively.

Simulating Pretraining. Given computational constraints, we cannot retrain BERT from scratch. Instead, we approximate the pretraining process by continuing to update BERT with MLM (Gururangan et al., 2020), which we refer to as second-stage pretraining. Although this is an approximation to the actual pretraining process, the second-stage pretraining shares the same fundamental problem for pretraining: how can unsupervised training lead to downstream performance gains?

We study the effectiveness of different masking strategies by comparing to a BERT model without second-stage pretraining (VANILLA). We experiment with three text classification datasets: SST-2 (Socher et al., 2013), Hyperpartisan (Kiesel et al., 2019), and AGNews (Zhang et al., 2015). SST-2 classifies movie reviews by binary sentiment; Hyperpartisan is a binary classification task on whether a news article takes an extreme partisan standpoint; and AGNews classifies news articles into four different topics. On SST-2 and AGNews, we perform the second-stage pretraining on the training inputs (not using the labels). On Hyperpartisan, we use 100k unlabeled news articles that are released with the dataset. For SST-2 and AGNews, we study a low-resource setting and set the number of finetuning examples to be 20. For Hyperpartisan, we use the training set, which has 515 labeled examples. All evaluations are performed by fine-tuning a `bert-base-uncased` model (See Appendix A for full details).

Approximating Cloze-like Masking. We cannot identify the optimal set of cloze-like masks for an arbitrary downstream task, but these three tasks have associated lexicons which we can use to approximate the cloze-like masks. For SST-2, we take the sentiment lexicon selected by Hu and Liu (2004); for Hyperpartisan, we take the NRC word-emotion association lexicon (Mohammad and Turney, 2013); and for AGNews, we extract topic words by training a logistic regression classifier and

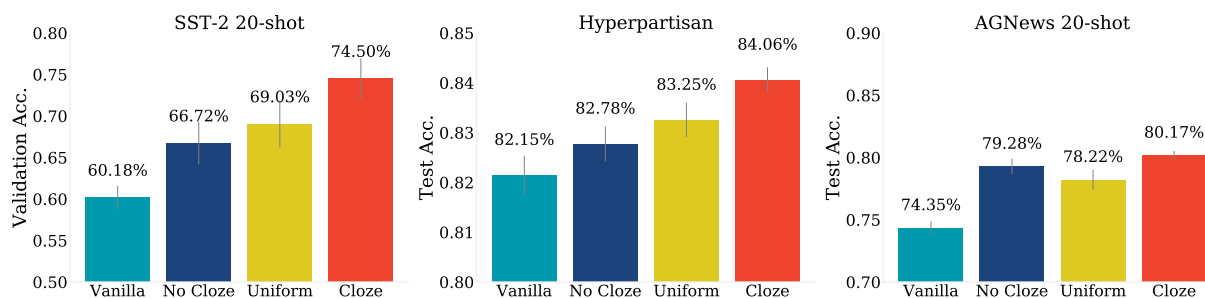


Figure 5: Finetuning performance with different masking strategies averaged across twenty random trials and error bars showing 95% confidence intervals. VANILLA represents a BERT model without any second-stage pretraining. CLOZE and NOCLOZE represent models train with or without cloze-like masks, respectively. UNIFORM uses the uniform random masking strategy proposed in Devlin et al. (2019) for second-stage pretraining.

taking the top 1k features to be cloze-like masks.

Results. Figure 5 plots the finetuning performance of different masking strategies. We observe that UNIFORM outperforms VANILLA, which indicates that second-stage pretraining is extracting useful information and our experiment setup is useful for studying how MLM leads to performance gains. As expected, CLOZE achieves the best accuracy, which confirms that cloze-like masks can be helpful and validates our cloze approximations.

The UNIFORM mask is much closer to NOCLOZE than CLOZE. This suggests that uniform masking does not leverage cloze-like masks well and cloze reductions alone cannot account for the success of MLM. This view is further supported by the observation that NOCLOZE outperforms VANILLA suggesting that generic masks that are not cloze-like still contain useful inductive biases.

Our results support our earlier view that there may be an alternative mechanism that allows generic masks that are not cloze-like to benefit downstream learning. Next, we will empirically examine BERT’s learned conditional independence structure among tokens and show that the statistical dependencies relate to syntactic dependencies.

5.2 Analysis: Unsupervised Parsing

Our analysis in section 4.1 shows that conditional MI (which is optimized by the MLM objective) can extract conditional independences. We will show that statistical dependencies estimated by conditional MI are related to syntactic dependencies by using conditional MI for unsupervised parsing.

Background. One might expect that the statistical dependencies among words are correlated with syntactic dependencies. Indeed, Futrell et al. (2019) show that heads and dependents in dependency parse trees have high pointwise mutual information (PMI) on average. However, previous at-

tempts (Carroll and Charniak, 1992; Paskin, 2002) show that unsupervised parsing approaches based on PMI achieve close to random accuracy. Our analysis suggests that MLMs extract a more fine-grained notion of statistical dependence (conditional MI) which does not suffer from the existence of latent variables (Proposition 3). We now show that the conditional MI captured by MLMs achieves far better performance, on par with classic unsupervised parsing baselines.

Baselines. We compare conditional MI to PMI as well as conditional PMI, an ablation in which we do not take expectation over possible words. For all statistical dependency based methods (cond. MI, PMI, and cond. PMI), we compute pairwise dependence for each word pair in a sentence and construct a minimum spanning tree on the negative values to generate parse trees. To contextualize our results, we compare against three simple baselines: RANDOM which draws a random tree on the input sentence, LINEARCHAIN which links adjacent words in a sentence, and a classic unsupervised parsing method (Klein and Manning, 2004).

Experimental Setup. We conduct experiments on the English Penn Treebank using the WSJ corpus and convert the annotated constituency parses to Stanford Dependency Formalism (de Marneffe et al., 2006). Following Yang et al. (2020), we evaluate on sentences of length ≤ 10 in the test split, which contains 389 sentences (Appendix B.1 describes the same experiment on longer sentences, which have similar results). We experiment with the `bert-base-cased` model (more details in Appendix A) and evaluate by the undirected unlabeled attachment score (UAS).

Results. Table 1 shows a much stronger-than-random association between conditional MI and dependency grammar. In fact, the parses extracted




Figure 6: An example parse extracted from conditional MI. The black parse tree above the sentence represents the ground-truth parse and the red parse below is extracted from conditional MI. The correctly predicted edges are labeled with the annotated relations, and the incorrect ones are labeled as wrong.

Method	UAS
RANDOM	28.50 ± 0.73
LINEARCHAIN	54.13
Klein and Manning (2004)	55.91 ± 0.68
PMI	33.94
CONDITIONAL PMI	52.44 ± 0.19
CONDITIONAL MI	58.74 ± 0.22

Table 1: Unlabeled Undirected Attachment Score on WSJ10 test split (section 23). Error bars show standard deviation across three random seeds.

from conditional MI has better quality than LINEARCHAIN and the classic method ([Klein and Manning, 2004](#)). Unlike conditional MI, PMI only has a close-to-random performance, which is consistent with prior work. We also see that conditional MI outperforms conditional PMI, which is consistent with our theoretical framework that suggests that conditional MI (and not PMI) recovers the graphical model structure.

We also perform a fine-grained analysis by investigating relations where conditional MI differs from LINEARCHAIN. Because the test split is small and conditional MI does not involve any training, we perform this analysis on 5,000 sentences from the training split. [Table 2](#) presents the results and shows that conditional MI does not simply recover the linear chain bias. Meanwhile, we also observe a deviation between conditional MI and dependency grammar on relations like `number` and `cc`. This is reasonable because certain aspects of dependency grammar depend on human conventions that do not necessarily have a consensus ([Popel et al., 2013](#)).

[Figure 6](#) illustrates with an example parse extracted from conditional MI. We observe that conditional MI correctly captures `dobj` and `conj`. Knowing the verb, e.g. *represents*, limits the range of objects that can appear in a sentence so intuitively we expect a high conditional MI between the direct object and the verb. Similarly for phrases like “*A and B*”, we would expect *A* and *B* to be statistically dependent. However, conditional MI fails

Relation	Conditional MI	Linear Chain
<code>xcomp</code>	48.18	9.93
<code>conj</code>	43.36	7.58
<code>dobj</code>	58.96	30.33
<code>number</code>	50.55	92.62
<code>quantmod</code>	56.82	72.73
<code>cc</code>	31.39	41.10

Table 2: Six relations on which conditional MI disagrees with LINEARCHAIN under log odds ratio test with $p = 0.05$. A comprehensive list is in [Appendix A](#).

to capture `cc` (between *apathy* and *or*). Instead, it links *or* with *either* which certainly has statistical dependence. This once again suggests that the ‘errors’ incurred by the conditional PMI method are not simply failures to estimate dependence but natural differences in the definition of dependence.

6 Discussion and Conclusion

We study how MLM with uniform masking can learn useful linguistic structures and inductive biases for downstream tasks. Our work demonstrates that a substantial part of the performance gains of MLM pretraining cannot be attributed to task-specific, cloze-like masks. Instead, learning with task-agnostic, generic masks encourages the model to capture direct statistical dependencies among tokens, and we show through unsupervised parsing evaluations that this has a close correspondence to syntactic structures. Existing work has suggested that statistical and syntactic dependencies are fundamentally different, with unsupervised parsing based on PMI achieving close-to-random performance. Our work demonstrates that this is not necessarily the case, and better measures of statistical dependence (such as those learned by MLMs) can serve as implicit supervision for learning syntactic structures. Our findings open new space for future works on how syntax can be learned in an emergent way and on how to design masking strategies that further improve dependency learning.

Acknowledgement

We thank Rishi Bommasani, Lisa Li, Kawin Ethayarajh, the anonymous reviewers and the Stanford NLP group for their helpful feedback and discussions.

References

- A. Anandkumar, V. Y. F. Tan, F. Huang, and A. S. Willsky. 2012. High-dimensional structure estimation in ising models: Local separation criterion. *Annals of Statistics*, 40(3):1346–1375.
- S. Arora, Y. Li, Y. Liang, T. Ma, and A. Risteski. 2015. Random walks on context spaces: Towards an explanation of the mysteries of semantic word embeddings. *arXiv preprint arXiv:1502.03520*.
- Y. Belinkov and J. Glass. 2019. Analysis methods in neural language processing: A survey. *Transactions of the Association for Computational Linguistics (ACL)*, 7:49–72.
- D. Blei, A. Ng, and M. I. Jordan. 2003. Latent Dirichlet allocation. *Journal of Machine Learning Research (JMLR)*, 3:993–1022.
- T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- G. Carroll and E. Charniak. 1992. Two experiments on learning probabilistic dependency grammars from corpora. In *AAAI Conference on Artificial Intelligence*.
- M. de Marneffe, B. MacCartney, and C. D. Manning. 2006. Generating typed dependency parses from phrase structure parses. In *LREC*.
- J. Devlin, M. Chang, K. Lee, and K. Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Association for Computational Linguistics (ACL)*, pages 4171–4186.
- R. Futrell, P. Qian, E. Gibson, E. Fedorenko, and I. Blank. 2019. Syntactic dependencies correspond to word pairs with high mutual information. In *Proceedings of the Fifth International Conference on Dependency Linguistics (Depling, SyntaxFest 2019)*, pages 3–13, Paris, France. Association for Computational Linguistics.
- S. Gururangan, A. Marasović, S. Swayamdipta, K. Lo, I. Beltagy, D. Downey, and N. A. Smith. 2020. Don’t stop pretraining: Adapt language models to domains and tasks. In *Association for Computational Linguistics (ACL)*, pages 8342–8360, Online. Association for Computational Linguistics.
- T. B. Hashimoto, D. Alvarez-Melis, and T. S. Jaakkola. 2016. Word embeddings as metric recovery in semantic spaces. *Transactions of the Association for Computational Linguistics (ACL)*, 4:273–286.
- J. He, G. Neubig, and T. Berg-Kirkpatrick. 2018. Unsupervised learning of syntactic structure with invertible neural projections. In *Empirical Methods in Natural Language Processing*, pages 1292–1302, Brussels, Belgium. Association for Computational Linguistics.
- J. Hewitt and C.D. Manning. 2019. A structural probe for finding syntax in word representations. In *North American Association for Computational Linguistics (NAACL)*, pages 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- M. Hu and B. Liu. 2004. Mining and summarizing customer reviews. In *International Conference on Knowledge Discovery and Data Mining (KDD)*, KDD ’04, page 168–177, New York, NY, USA. Association for Computing Machinery.
- Z. Jiang, F. F. Xu, J. Araki, and G. Neubig. 2020. How can we know what language models know? *Transactions of the Association for Computational Linguistics (ACL)*, 8:423–438.
- J. Kiesel, M. Mestre, R. Shukla, E. Vincent, P. Adineh, D. Corney, B. Stein, and M. Potthast. 2019. SemEval-2019 task 4: Hyperpartisan news detection. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 829–839, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- D. Kingma and J. Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- D. Klein and C.D. Manning. 2004. Corpus-based induction of syntactic structure: Models of dependency and constituency. In *Association for Computational Linguistics (ACL)*, pages 478–485, Barcelona, Spain.
- J. D. Lee, Q. Lei, N. Saunshi, and J. Zhuo. 2020. Predicting what you already know helps: Provable self-supervised learning. *arXiv preprint arXiv:2008.01064*.
- O. Levy and Y. Goldberg. 2014. Neural word embedding as implicit matrix factorization. In *Advances in Neural Information Processing Systems*, volume 27, pages 2177–2185. Curran Associates, Inc.
- N. F. Liu, M. Gardner, Y. Belinkov, M. E. Peters, and N. A. Smith. 2019a. Linguistic knowledge and transferability of contextual representations. In *North American Association for Computational Linguistics (NAACL)*, pages 1073–1094, Minneapolis, Minnesota. Association for Computational Linguistics.

- Y. Liu. 2019. Fine-tune BERT for extractive summarization. *arXiv preprint arXiv:1903.10318*.
- Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov. 2019b. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*.
- C. D. Manning, M. Surdeanu, J. Bauer, J. Finkel, S. J. Bethard, and D. McClosky. 2014. The stanford coreNLP natural language processing toolkit. In *ACL system demonstrations*.
- N. Meinshausen and P. Bühlmann. 2006. High-dimensional graphs and variable selection with the lasso. *Annals of Statistics*, 34(3):1436–1462.
- T. Mikolov, K. Chen, G. Corrado, and Jeffrey. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- S. M. Mohammad and P. D. Turney. 2013. Crowdsourcing a word-emotion association lexicon. *Computational Intelligence*, 29(3):436–465.
- B. Pang, L. Lee, and S. Vaithyanathan. 2002. Thumbs up? sentiment classification using machine learning techniques. In *Empirical Methods in Natural Language Processing*, pages 79–86. Association for Computational Linguistics.
- M. A. Paskin. 2002. Grammatical bigrams. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- J. Pennington, R. Socher, and C. D. Manning. 2014. GloVe: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- M. Peters, M. Neumann, L. Zettlemoyer, and W. Yih. 2018. Dissecting contextual word embeddings: Architecture and representation. In *Empirical Methods in Natural Language Processing*, pages 1499–1509, Brussels, Belgium. Association for Computational Linguistics.
- F. Petroni, T. Rocktäschel, S. Riedel, P. Lewis, A. Bakhtin, Y. Wu, and A. Miller. 2019. Language models as knowledge bases? In *Empirical Methods in Natural Language Processing and International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2463–2473, Hong Kong, China. Association for Computational Linguistics.
- M. Popel, D. Mareček, J. Štěpánek, D. Zeman, and Z. Žabokrtský. 2013. Coordination structures in dependency treebanks. In *Association for Computational Linguistics (ACL)*, pages 517–527, Sofia, Bulgaria. Association for Computational Linguistics.
- A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8).
- N. Saunshi, S. Malladi, and S. Arora. 2020. A mathematical exploration of why language models help solve downstream tasks. *arXiv preprint arXiv:2010.03648*.
- R. Socher, A. Perelygin, J. Y. Wu, J. Chuang, C. D. Manning, A. Y. Ng, and C. Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- G. W. Stewart and J. Sun. 1990. *Matrix Perturbation Theory*. Academic Press.
- Alex Tamkin, Trisha Singh, Davide Giovanardi, and Noah Goodman. 2020. Investigating transferability in pretrained language models. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1393–1401. Association for Computational Linguistics.
- I. Tenney, P. Xia, B. Chen, A. Wang, A. Poliak, R. T. McCoy, N. Kim, B. Van Durme, S. Bowman, D. Das, and E. Pavlick. 2019. What do you learn from context? probing for sentence structure in contextualized word representations. In *International Conference on Learning Representations (ICLR)*.
- R. Tibshirani. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288.
- D. Wadden, U Wennberg, Y. Luan, and H. Hajishirzi. 2019. Entity, relation, and event extraction with contextualized span representations. In *Empirical Methods in Natural Language Processing and International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5784–5789, Hong Kong, China. Association for Computational Linguistics.
- A. Wang and K. Cho. 2019. BERT has a mouth, and it must speak: BERT as a Markov random field language model. *arXiv preprint arXiv:1902.04094*.
- Z. Wu, Y. Chen, B. Kao, and Q. Liu. 2020. Perturbed masking: Parameter-free probing for analyzing and interpreting BERT. In *Association for Computational Linguistics (ACL)*, pages 4166–4176, Online. Association for Computational Linguistics.
- S. Yang, Y. Jiang, W. Han, and K. Tu. 2020. Second-order unsupervised neural dependency parsing. *arXiv preprint arXiv:2010.14720*.
- T. Zhang, F. Wu, A. Katiyar, K. Q. Weinberger, and Y. Artzi. 2020. Revisiting few-sample bert fine-tuning. *arXiv preprint arXiv:2006.05987*.
- X. Zhang, J. Zhao, and Y. LeCun. 2015. Character-level convolutional networks for text classification. In *Advances in Neural Information Processing Systems*, volume 28, pages 649–657. Curran Associates, Inc.

J. Zhu, Y. Xia, L. Wu, D. He, T. Qin, W. Zhou, H. Li, and T. Liu. 2020. Incorporating BERT into neural machine translation. In *International Conference on Learning Representations (ICLR)*.

Dataset	# Classes	# Pretrain	# Finetune	# Test
SST-2	2	67k	20	1.8k
Hyperpartisan	2	100k	515	130
AGNews	4	113k	20	6.7k

Table 3: Specifications of datasets. For AGNews, we put away 6.7k as a development set.

A Experimental Details

Experimental details for Section 3.2 Our transformers have 2 layers and for each transformer block, the hidden size and the intermediate size are both 64. We finetune the models for 10 epochs and apply early stopping based on validation accuracy. We use Adam (Kingma and Ba, 2014) for optimization, using a learning rate of $1e^{-3}$ for pretraining and $1e^{-4}$ for finetuning.

Experimental details for Section 5.1 Table 3 summarizes the dataset statistics of three real-world datasets we studied. For second stage pretraining, we update the BERT model for 10 epochs. Following the suggestion in Zhang et al. (2020), we finetune the pretrained BERT models for 400 steps, using a batch size of 16 and a learning rate of $1e^{-5}$. We apply linear learning rate warmup for the first 10% of finetuning and linear learning rate decay for the rest. For SST-2 and AGNews, we average the results over 20 random trials. For Hyperpartisan, because the test set is small and the variation is larger, we average the results over 50 random trials and evaluate on the union the development set and the test set for more stable results.

Experimental details for Section 5.2 We convert the annotated constituency parses using the Stanford CoreNLP package (Manning et al., 2014). We compute conditional MI and conditional PMI using the bert-base-cased model and run Gibbs sampling for 2000 steps. BERT’s tokenization may split a word into multiple word pieces. We aggregate the dependencies between a word and multiple word pieces by taking the maximum value. We compute the PMI statistics and train the K&M model (Klein and Manning, 2004) on sentences of length ≤ 10 in the WSJ train split (section 2-21). For DMV, we train with the annotated POS tags using a public implementation released by (He et al., 2018). Results are averaged over three runs when applicable.

B Additional Results

B.1 Additional Results in Section 5.2

We conduct an additional experiment on the English Penn Treebank to verify that conditional MI can extract parses for sentences longer than ten words. To expedite experimentation, we subsample 200 out of 2416 sentences from the test split of English Penn Treebank and the average sentence length of our subsampled dataset is 24.1 words. When applicable, we average over three random seeds and report standard deviations. Table 4 presents the UUAS of conditional MI and other methods. We draw similar conclusions as in Section 5.2, observing that the parses drawn by conditional MI have higher quality than those of other baselines.

Table 5 presents a comprehensive list of relations on which Conditional MI disagrees with LINEARCHINA under a log odds ratio test with $p = 0.05$.

Method	UUAS
RANDOM	9.14 ± 0.42
LINEARCHAIN	47.69
Klein and Manning (2004)	48.76 ± 0.24
PMI	28.05
CONDITIONAL PMI	44.75 ± 0.09
CONDITIONAL MI	50.62 ± 0.38

Table 4: Unlabeled Undirected Attachment Score on subsampled WSJ test split (section 23). Error bars show standard deviation across three random seeds.

Relation	Conditional MI	Linear Chain
xcomp	48.18	9.93
conj	43.36	7.58
nsubjpass	33.81	0.47
dobj	58.96	30.33
mark	30.71	9.45
poss	58.63	40.96
ccomp	20.92	4.18
vmod	55.32	41.84
tmod	39.25	27.68
dep	50.15	40.03
pobj	48.68	40.79
nsub	55.87	48.69
number	50.55	92.62
possessive	72.00	97.78
pcomp	60.00	77.00
quantmod	56.82	72.73
appos	55.56	70.59
num	65.11	76.49
cc	31.39	41.10
prep	56.41	66.12
auxpass	75.00	83.26
nn	72.97	77.88
aux	55.49	59.66

Table 5: All relations on which Conditional MI disagree with LINEARCHAIN under a log odds ratio test with $p = 0.05$.

C Proofs

Proof of Proposition 2 We first recall our statement.

Proposition 2. *Let λ_k be the k th eigenvalue of $A\Sigma_{\mathbf{Z}\mathbf{Z}}A^\top$ and $\lambda_{\mathbf{X}\mathbf{X},k+1}$ be the $k+1$ th eigenvalue of $\Sigma_{\mathbf{X}\mathbf{X}}$ and V be the first k eigenvectors of $\text{Cov}(\mathbf{X})$. Assuming $\lambda_k > \lambda_{\mathbf{X}\mathbf{X},k+1}$, we have*

$$\mathbb{E}_{\mathbf{X}} \|\mathbf{AZ} - \mathbf{X}_{\text{PCA}}\|_2 \leq \frac{\sqrt{2} \|\Sigma_{\mathbf{X}\mathbf{X}}\|_{op}}{\lambda_k - \lambda_{\mathbf{X}\mathbf{X},k+1}} (\|\mathbf{AZ}\|_2 + \sqrt{\text{tr}(\Sigma_{\mathbf{X}\mathbf{X}})}) + \|AA^\top\|_{op} \sqrt{\text{tr}(\Sigma_{\mathbf{X}\mathbf{X}})},$$

where $\|\cdot\|_{op}$ is the operator norm and $\text{tr}(\cdot)$ is the trace.

Proof

We will use the Davis-Kahan Theorem for our proof.

Theorem (Davis-Kahan (Stewart and Sun, 1990)). *Let σ be the eigengap between the k th and the $k+1$ th eigenvalue of two positive semidefinite symmetric matrices Σ and Σ' . Also, let V and V' be the first k eigenvectors of Σ and Σ' respectively. Then we have,*

$$\frac{1}{\sqrt{2}} \|VV^\top - V'V'^\top\|_{op} \leq \frac{\|\Sigma - \Sigma'\|_{op}}{\sigma}.$$

That is, we can bound the error in the subspace projection in terms of the matrix perturbation.

In our setting, we choose $\Sigma = A\Sigma_{\mathbf{Z}\mathbf{Z}}A^\top + \Sigma_{\mathbf{X}\mathbf{X}}$ and $\Sigma' = A\Sigma_{\mathbf{Z}\mathbf{Z}}A^\top$. We know the eigengap of Σ' is λ_k because Σ' only has k nonzero eigenvalues. By Weyl's inequality, the k th eigenvalue is at most perturbed by $\lambda_{\mathbf{X}\mathbf{X},k+1}$, which is the $k+1$ eigenvalue of $\Sigma_{\mathbf{X}\mathbf{X}}$. Let V be the top k eigenvectors of Σ' and assuming $\lambda_k > \lambda_{\mathbf{X}\mathbf{X},k+1}$, we have,

$$\begin{aligned} \frac{1}{\sqrt{2}} \|AA^\top - VV^\top\|_{op} &\leq \frac{\|\Sigma - \Sigma'\|_{op}}{\lambda_k - \lambda_{\mathbf{X}\mathbf{X},k+1}} \\ &= \frac{\|\Sigma_{\mathbf{X}\mathbf{X}}\|_{op}}{\lambda_k - \lambda_{\mathbf{X}\mathbf{X},k+1}}. \end{aligned}$$

Turning this operator norm bound into approximation bound, we have

$$\begin{aligned} \mathbb{E}_{\mathbf{X}} \|\mathbf{AZ} - \mathbf{X}_{\text{PCA}}\|_2 &= \mathbb{E}_{\mathbf{X}} \|AA^\top \mathbf{AZ} - VV^\top \mathbf{X}\|_2 \\ &= \mathbb{E}_{\mathbf{X}} \|AA^\top \mathbf{AZ} - VV^\top \mathbf{AZ} + VV^\top \mathbf{AZ} - VV^\top \mathbf{X}\|_2 \\ &\leq \mathbb{E}_{\mathbf{X}} \|AA^\top \mathbf{AZ} - VV^\top \mathbf{AZ}\|_2 + \|VV^\top \mathbf{AZ} - VV^\top \mathbf{X}\|_2 \\ &\leq \mathbb{E}_{\mathbf{X}} \|AA^\top \mathbf{AZ} - VV^\top \mathbf{AZ}\|_2 + \|VV^\top (\mathbf{AZ} - \mathbf{X})\|_2 \\ &\leq \mathbb{E}_{\mathbf{X}} \|AA^\top - VV^\top\|_{op} \cdot \|\mathbf{AZ}\|_2 + \|VV^\top\|_{op} \|\mathbf{AZ} - \mathbf{X}\|_2 \\ &= \mathbb{E}_{\mathbf{X}} \|AA^\top - VV^\top\|_{op} \cdot \|\mathbf{AZ}\|_2 + \|AA^\top + VV^\top - AA\|_{op} \|\mathbf{AZ} - \mathbf{X}\|_2 \\ &\leq \mathbb{E}_{\mathbf{X}} \|AA^\top - VV^\top\|_{op} \cdot \|\mathbf{AZ}\|_2 + (\|AA^\top\|_{op} + \|VV^\top - AA\|_{op}) \|\mathbf{AZ} - \mathbf{X}\|_2 \\ &= \mathbb{E}_{\mathbf{X}} \|AA^\top - VV^\top\|_{op} \cdot (\|\mathbf{AZ}\|_2 + \|\mathbf{AZ} - \mathbf{X}\|_2) + \|AA^\top\|_{op} \|\mathbf{AZ} - \mathbf{X}\|_2. \end{aligned}$$

We use the fact that $\mathbb{E}_{\mathbf{X},\mathbf{Z}} \|\mathbf{AZ} - \mathbf{X}\|_2^2 = \text{tr}(\Sigma_{\mathbf{X}\mathbf{X}})$ and Jensen's inequality to bound,

$$\mathbb{E}_{\mathbf{X}} \|\mathbf{AZ} - \mathbf{X}\|_2 \leq \sqrt{\text{tr}(\Sigma_{\mathbf{X}\mathbf{X}})}.$$

Combining these inequalities, we have

$$\begin{aligned} & \mathbb{E}_{\mathbf{X}} \|\mathbf{AZ} - \mathbf{X}_{\text{PCA}}\|_2 \\ & \leq \frac{\sqrt{2} \|\Sigma_{\mathbf{X}\mathbf{X}}\|_{op}}{\lambda_k - \lambda_{\mathbf{X}\mathbf{X},k+1}} \cdot (\|\mathbf{AZ}\|_2 + \sqrt{\text{tr}(\Sigma_{\mathbf{X}\mathbf{X}})}) + \left\| \mathbf{AA}^\top \right\|_{op} \sqrt{\text{tr}(\Sigma_{\mathbf{X}\mathbf{X}})} \end{aligned}$$

□

Proof of Proposition 1 We first recall our statement.

Proposition 1. *Assuming that $\Sigma_{\mathbf{X}\mathbf{X}}$ is full rank,*

$$\mathbf{x}_{\text{mask},i} = \beta_{2\text{SLS},i} \mathbf{X}_{\setminus i} + O(\|\Sigma_{\mathbf{X}\mathbf{X},\setminus i,i}\|_2),$$

Proof Let $A_{\setminus i} \in \mathbb{R}^{d-1 \times k}$ be the matrix where we omit the i th row of A and $A_i \in \mathbb{R}^k$ be the i th row of A . Let $\Sigma_{\mathbf{X}\mathbf{X},\setminus i,\setminus i} \in \mathbb{R}^{d-1 \times d-1}$ be the matrix where we omit the i th row and i th column of $\Sigma_{\mathbf{X}\mathbf{X}}$, and $\Sigma_{\mathbf{X}\mathbf{X},\setminus i,i} \in \mathbb{R}^{d-1}$ be the vector formed by dropping the i th row and taking the i th column of $\Sigma_{\mathbf{X}\mathbf{X}}$. Similarly, denote $\mathbf{X}_{\setminus i} \in \mathbb{R}^{d-1}$ be the vector where we omit the i coordinate of \mathbf{X} .

We start by writing down the expression of $\beta_{2\text{SLS},i}$. Recall that the Least Squares regression between two zero-mean Gaussian variables \mathbf{X} and \mathbf{Y} can be written as

$$\beta = \text{Cov}(\mathbf{X}, \mathbf{Y}) \text{Cov}(\mathbf{X}, \mathbf{X})^{-1},$$

where $\text{Cov}(\mathbf{X}, \mathbf{X})$ is the covariance matrix of X and we assume it is full rank. Since $\text{Cov}(\mathbf{X}_{\setminus i}, \mathbf{Z})$ is $A_{\setminus i} \Sigma_{\mathbf{Z}\mathbf{Z}}$, we can write the coefficient of regression from $\mathbf{X}_{\setminus i}$ to \mathbf{Z} as

$$\beta_{\mathbf{X}_{\setminus i} \rightarrow \mathbf{Z}} = \Sigma_{\mathbf{Z}\mathbf{Z}} A_{\setminus i}^\top (A_{\setminus i} \Sigma_{\mathbf{Z}\mathbf{Z}} A_{\setminus i}^\top + \Sigma_{\mathbf{X}\mathbf{X},\setminus i,\setminus i})^{-1}$$

and by assumption we have $\beta_{\mathbf{Z} \rightarrow \mathbf{x}_i} = A_i$. So we can write down

$$\beta_{2\text{SLS},i} = A_i \Sigma_{\mathbf{Z}\mathbf{Z}} A_{\setminus i}^\top (A_{\setminus i} \Sigma_{\mathbf{Z}\mathbf{Z}} A_{\setminus i}^\top + \Sigma_{\mathbf{X}\mathbf{X},\setminus i,\setminus i})^{-1}.$$

Now we consider masked regression for the i th coordinate, \mathbf{x}_i ,

$$\beta_{\mathbf{X}_{\setminus i} \rightarrow \mathbf{x}_i} = (A_i \Sigma_{\mathbf{Z}\mathbf{Z}} A_{\setminus i}^\top + \Sigma_{\mathbf{X}\mathbf{X},\setminus i,i}) (A_{\setminus i} \Sigma_{\mathbf{Z}\mathbf{Z}} A_{\setminus i}^\top + \Sigma_{\mathbf{X}\mathbf{X},\setminus i,\setminus i})^{-1}.$$

Comparing $\beta_{2\text{SLS}}$ and $\beta_{\mathbf{X}_{\setminus i} \rightarrow \mathbf{x}_i}$, we observe that the second term is the same and the key is to bound the first term. Consider the error term between the coefficients,

$$\begin{aligned} & \left\| \Sigma_{\mathbf{X}\mathbf{X},\setminus i,i} (A_{\setminus i} \Sigma_{\mathbf{Z}\mathbf{Z}} A_{\setminus i}^\top + \Sigma_{\mathbf{X}\mathbf{X},\setminus i,\setminus i})^{-1} \right\|_2 \\ & \leq \left\| \Sigma_{\mathbf{X}\mathbf{X},\setminus i,i} \right\|_2 \left\| (A_{\setminus i} \Sigma_{\mathbf{Z}\mathbf{Z}} A_{\setminus i}^\top + \Sigma_{\mathbf{X}\mathbf{X},\setminus i,\setminus i})^{-1} \right\|_{op} \\ & \leq \left\| \Sigma_{\mathbf{X}\mathbf{X},\setminus i,i} \right\|_2 \left\| (A \Sigma_{\mathbf{Z}\mathbf{Z}} A^\top + \Sigma_{\mathbf{X}\mathbf{X}})^{-1} \right\|_{op}. \end{aligned}$$

That is, the error term scales with the off-diagonal terms $\|\Sigma_{\mathbf{X}\mathbf{X},\setminus i,i}\|_2$.

Converting our bound on the error term into an approximation bound, we have

$$\mathbf{x}_{\text{mask},i} = \beta_{2\text{SLS},i} \mathbf{X} + O(\|\Sigma_{\mathbf{X}\mathbf{X},\setminus i,i}\|_2).$$

□

Proof for Proposition 3.

Proposition 3. *The gap between conditional MI with and without latent variables is bounded by the conditional entropy $H(\mathbf{Z}|X_{\setminus\{i,j\}})$,*

$$I(x_i; x_j | X_{\setminus\{i,j\}}) - I(x_i; x_j | \mathbf{Z}, X_{\setminus\{i,j\}}) \leq 2H(\mathbf{Z}|X_{\setminus\{i,j\}}).$$

Proof The proof follows from the definition of conditional mutual information. Denote $H(\cdot)$ as the entropy function.

We start by observing that

$$\begin{aligned} I(x_i; x_j | \mathbf{Z}, X_{\setminus\{i,j\}}) &= I(x_i; x_j | X_{\setminus\{i,j\}}) - I(x_i; \mathbf{Z} | X_{\setminus\{i,j\}}) + I(x_i; \mathbf{Z} | x_j, X_{\setminus\{i,j\}}) \\ &\quad \text{(Through chain rule of mutual information.)} \\ &= I(x_i; x_j | X_{\setminus\{i,j\}}) + H(\mathbf{Z} | x_i, X_{\setminus\{i,j\}}) - H(\mathbf{Z} | X_{\setminus\{i,j\}}) \\ &\quad + H(\mathbf{Z} | x_j, X_{\setminus\{i,j\}}) - H(\mathbf{Z} | x_i, x_j, X_{\setminus\{i,j\}}). \end{aligned}$$

Then we have,

$$\begin{aligned} &I(x_i; x_j | X_{\setminus\{i,j\}}) - I(x_i; x_j | \mathbf{Z}, X_{\setminus\{i,j\}}) \\ &= -H(\mathbf{Z} | x_i, X_{\setminus\{i,j\}}) + H(\mathbf{Z} | X_{\setminus\{i,j\}}) - H(\mathbf{Z} | x_j, X_{\setminus\{i,j\}}) + H(\mathbf{Z} | x_i, x_j, X_{\setminus\{i,j\}}) \\ &\leq H(\mathbf{Z} | X_{\setminus\{i,j\}}) + H(\mathbf{Z} | x_i, x_j, X_{\setminus\{i,j\}}) \\ &\leq 2 \cdot H(\mathbf{Z} | X_{\setminus\{i,j\}}). \end{aligned}$$

□

Proposition 4. *Let*

$$\hat{I}_{p_\theta} = \mathbb{E}_{x_i, x_j} [\log p_\theta(x_i | X_{\setminus i}(j, x_j)) - \log \mathbb{E}_{x_j | x_i} p_\theta(x_i | X_{\setminus i}(j, x_j))]$$

be an estimator constructed by the model distribution p_θ . Then we can show,

$$|\hat{I}_{p_\theta} - I_p| \leq \mathbb{E}_{x_j} D_{\text{kl}}(p(x_i | X_{\setminus i}(j, x_j)) \| p_\theta(x_i | X_{\setminus i}(j, x_j))),$$

where D_{kl} represents the KL-divergence.

Proof Expanding the definition of mutual information, we write

$$\begin{aligned} I(x_i; x_j | X_{\setminus\{i,j\}}) - \hat{I}_\theta(x_i; x_j | X_{\setminus\{i,j\}}) &= \mathbb{E}_{x_j} [D_{\text{kl}}(p(x_i | x_j, X_{\setminus\{i,j\}}) \| p_\theta(x_i | x_j, X_{\setminus\{i,j\}}))] - \\ &\quad D_{\text{kl}}(\mathbb{E}_{x_j} p(x_i | x_j, X_{\setminus\{i,j\}}) \| \mathbb{E}_{x_j} p_\theta(x_i | x_j, X_{\setminus\{i,j\}})). \end{aligned}$$

Dropping the the second term, we have

$$\hat{I}_\theta(x_i; x_j | X_{\setminus\{i,j\}}) - I(x_i; x_j | X_{\setminus\{i,j\}}) \geq -\mathbb{E}_{x_j} [D_{\text{kl}}(p(x_i | x_j, X_{\setminus\{i,j\}}) \| p_\theta(x_i | x_j, X_{\setminus\{i,j\}}))].$$

Dropping the the first term, we have

$$\begin{aligned} &I(x_i; x_j | X_{\setminus\{i,j\}}) - \hat{I}_\theta(x_i; x_j | X_{\setminus\{i,j\}}) \\ &\leq D_{\text{kl}}(\mathbb{E}_{x_j} p(x_i | x_j, X_{\setminus\{i,j\}}) \| \mathbb{E}_{x_j} p_\theta(x_i | x_j, X_{\setminus\{i,j\}})) \\ &\leq \mathbb{E}_{x_j} D_{\text{kl}}(p(x_i | x_j, X_{\setminus\{i,j\}}) \| p_\theta(x_i | x_j, X_{\setminus\{i,j\}})), \end{aligned}$$

which uses the convexity of KL-divergence and Jensen's inequality.

□